



# Outline

- What is AmLight?
- Network Connectivity
- New AmLight SDN Architecture: Plane by Plane
- What have we planned for SC23?

# What is AmLight?

- A distributed academic exchange point built to enable collaboration among Latin America, Caribbean, Africa, and the U.S.
- Supported by NSF via the IRNC program under award # OAC-2029283 for the 2021-2025.
- Partnerships with R&E networks in the U.S., Latin America, Caribbean and Africa.



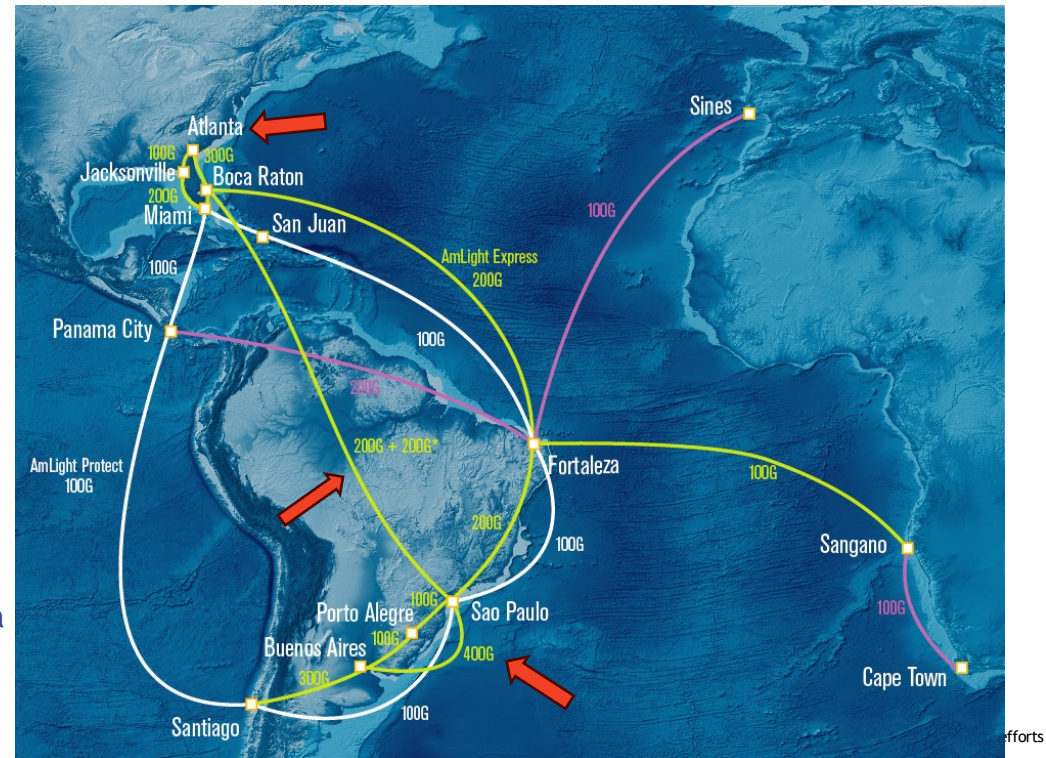


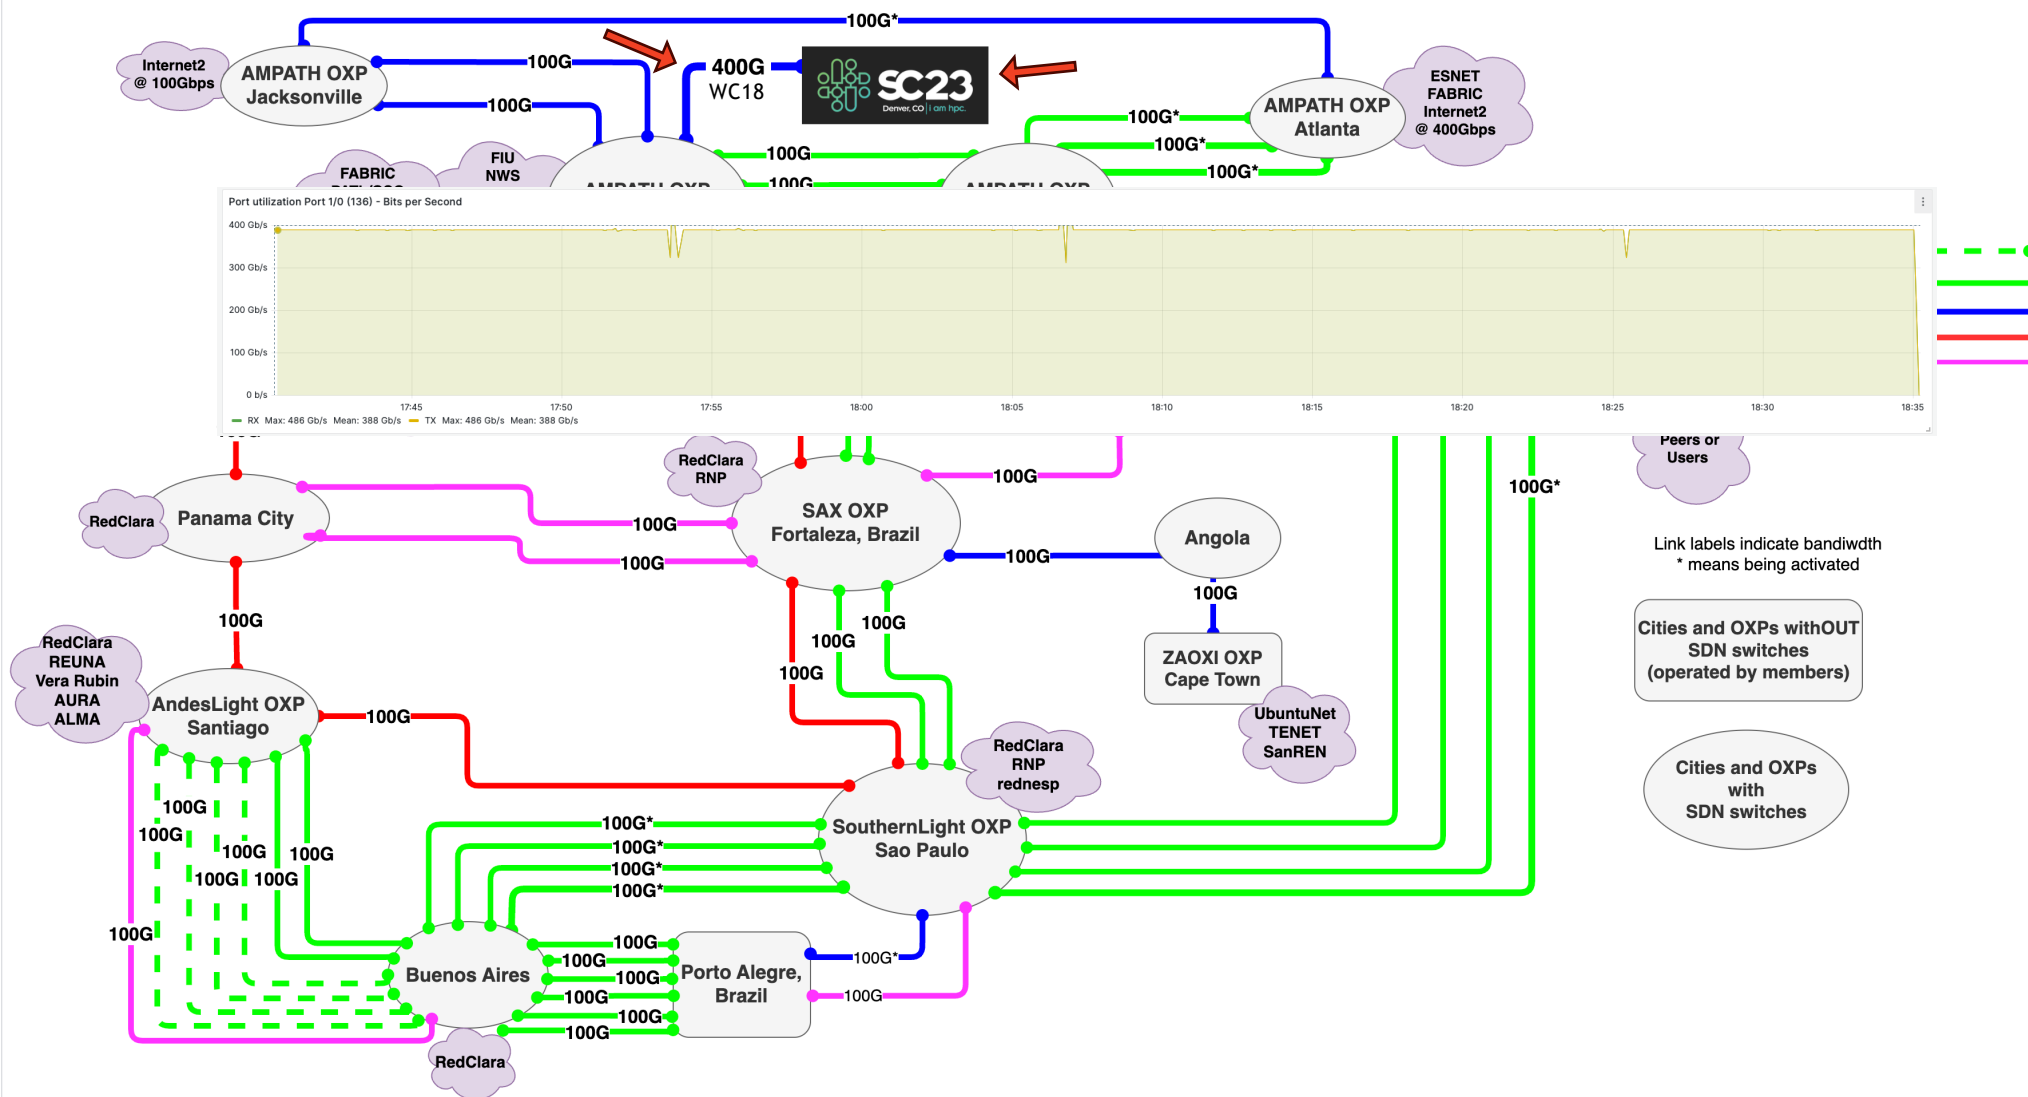
Network Connectivity...



# Network Connectivity

- 2.1+ Tbps of international connectivity
  - 600Gbps of upstream capacity between the U.S. and Latin America, and 100Gbps to Africa
  - +400Gbps in 2023 and +200Gbps in 2024
- By 2025, AmLight will reach 4.9 Tbps of total capacity
- NAPs:
  - Florida (Miami, Boca Raton, Jacksonville), Brazil (Sao Paulo, Fortaleza), Chile (Santiago), Puerto Rico (San Juan), Panama (Panama City), South Africa (Cape Town),
  - New: Georgia (Atlanta), Argentina (Buenos Aires)



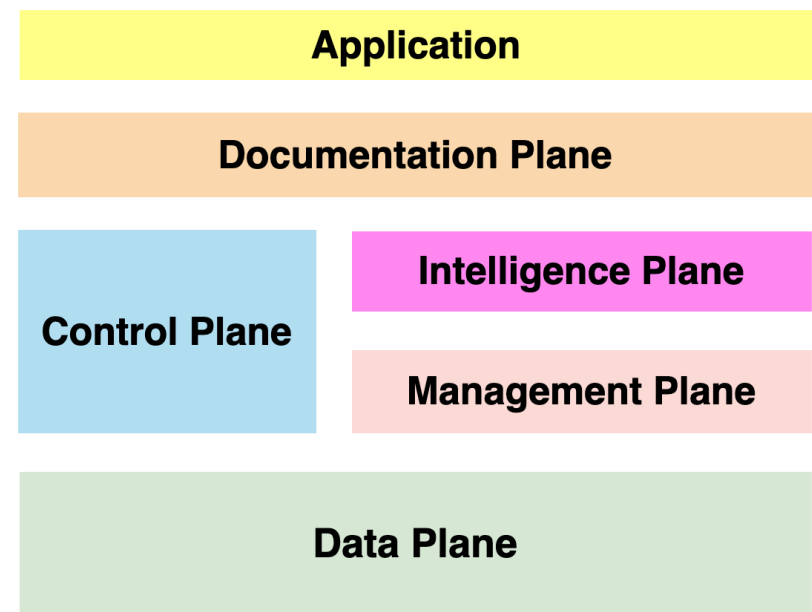


# AmLight SDN Architecture...

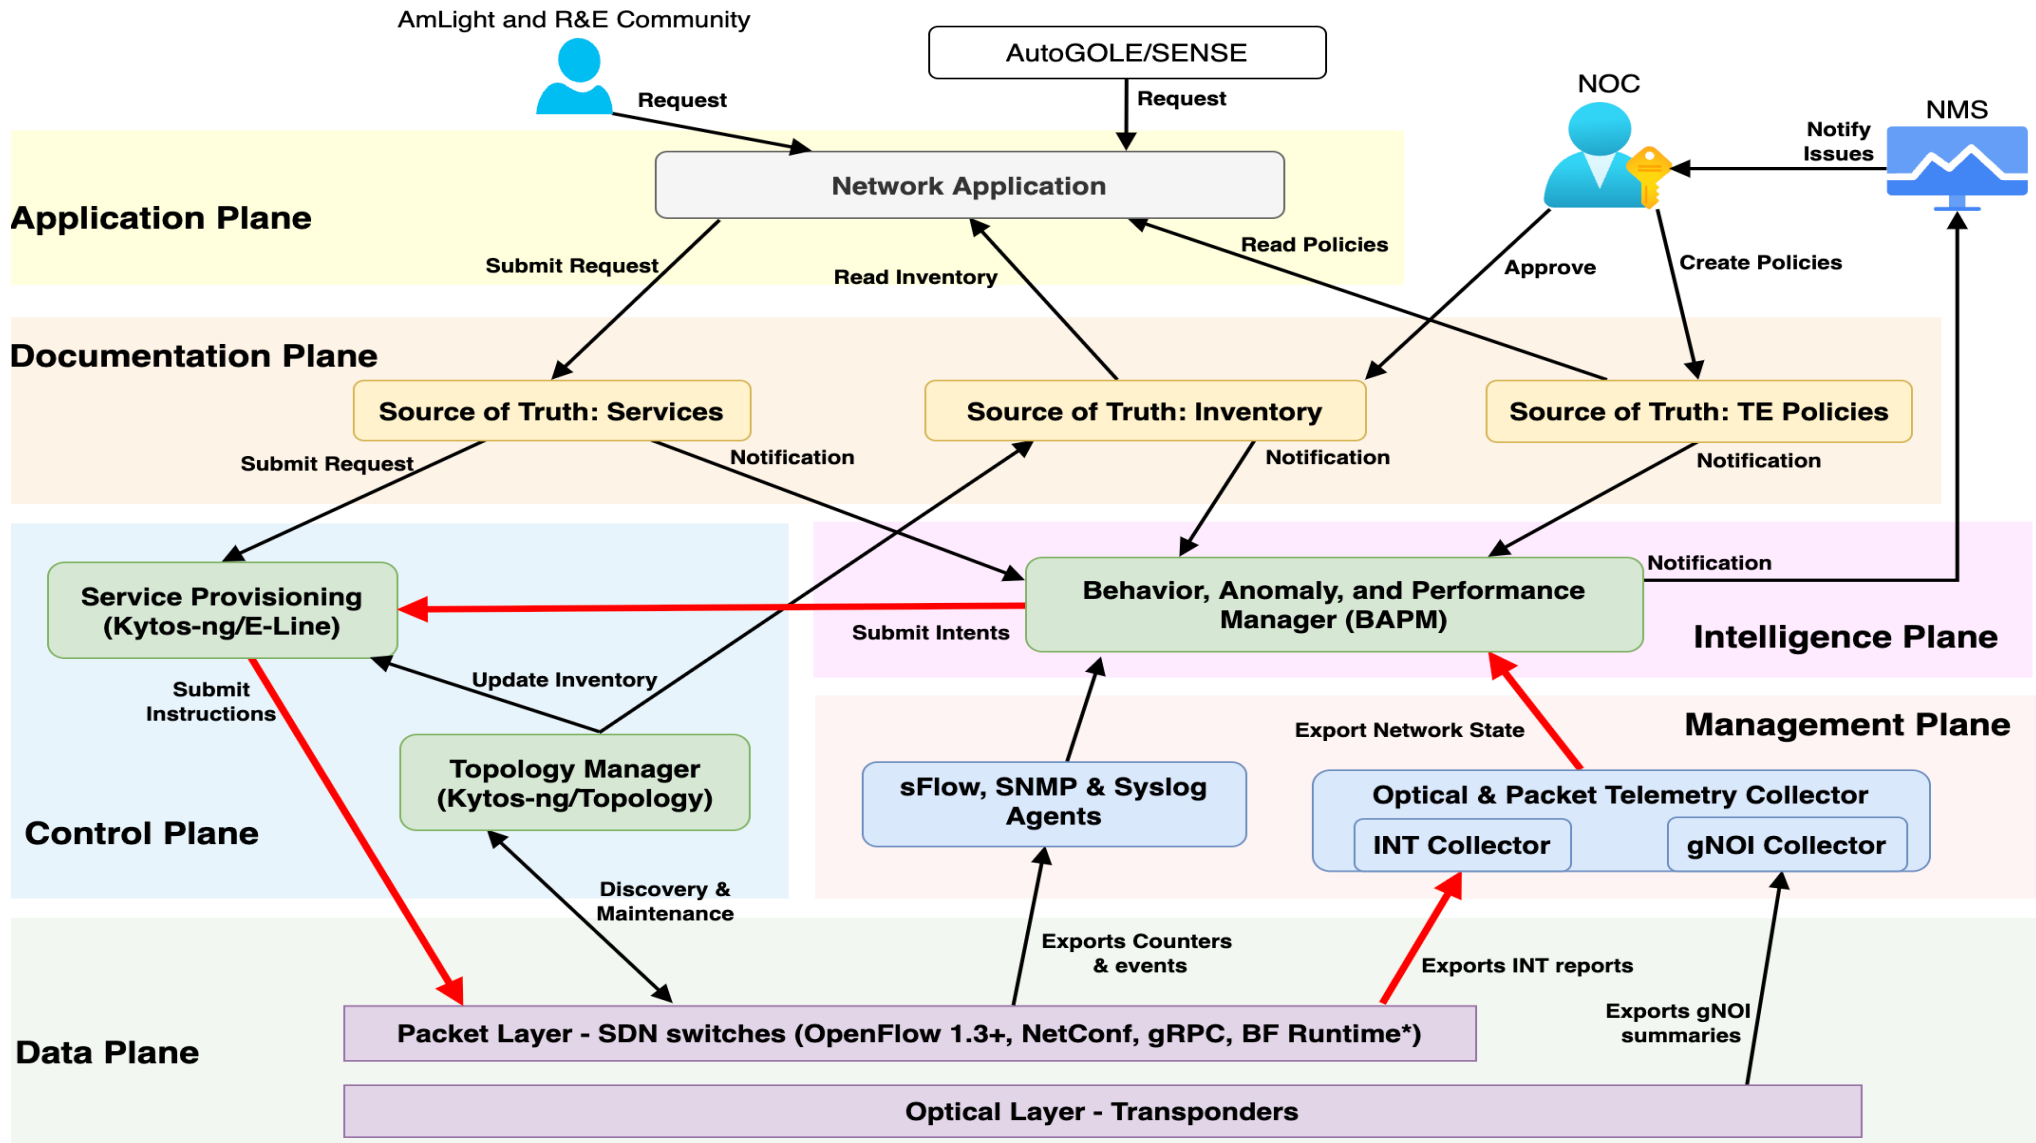
# AmLight SDN Architecture - 2021-2025

For 2021-2025, we will add specialized components per SDN Plane:

- New Control Plane:
  - Kytos-ng SDN Platform developed to address AmLight's and science drivers' requirements
- New Data Plane:
  - Programmable switches replacing legacy devices with customizable P4 pipeline
- New Management Plane:
  - Granular real-time network visibility with In-band Network Telemetry (INT)
- New Intelligence Plane:
  - Traffic engineering and optimizations based on granular telemetry reports in real-time

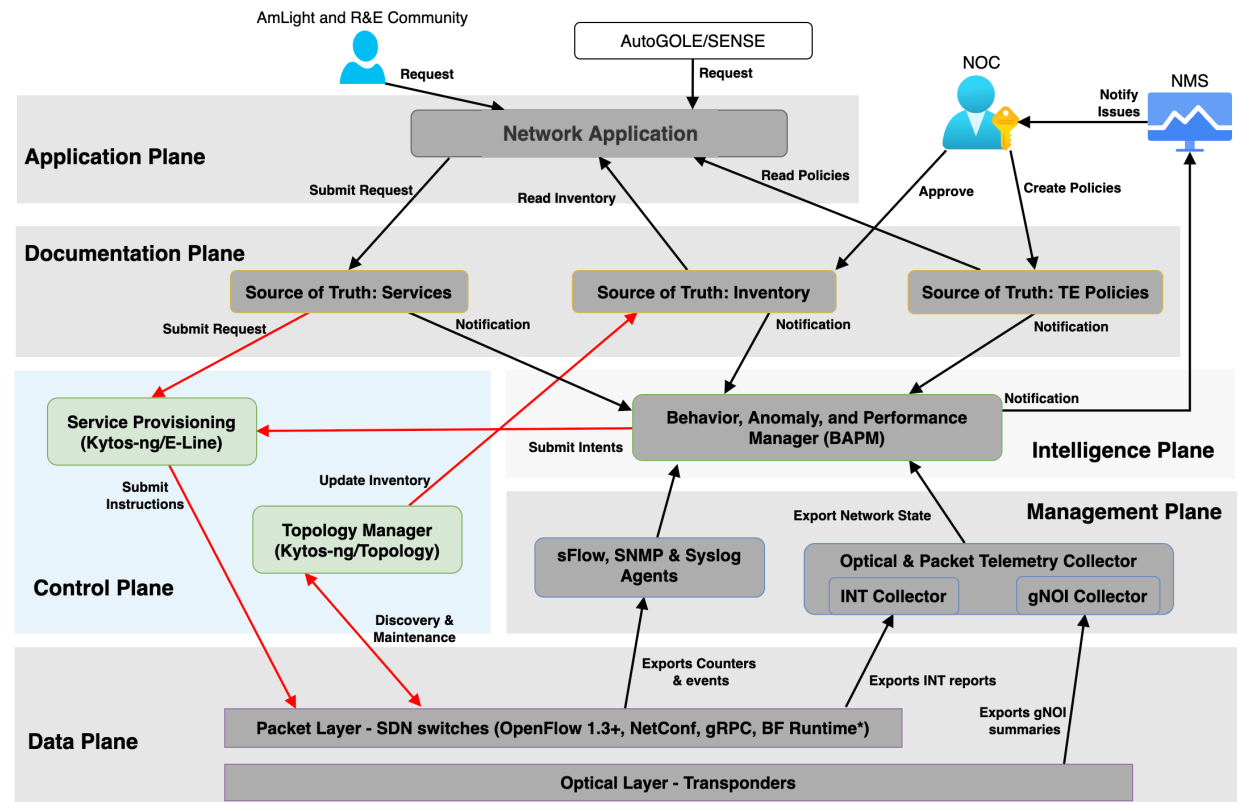






# Control Plane

- **Major transformation** compared to the previous AmLight-Exp project
  - From multiple open-source solutions (ONOS, OESS, FloodLight, and Ryu) to a new solution (Kytos-ng)
- **Brand new SDN controller built from scratch**
- Built to be fully compatible with the new AmLight Data Plane
- Addresses the requirements of the AmLight community and science drivers:
  - Telemetry and pathfinding options



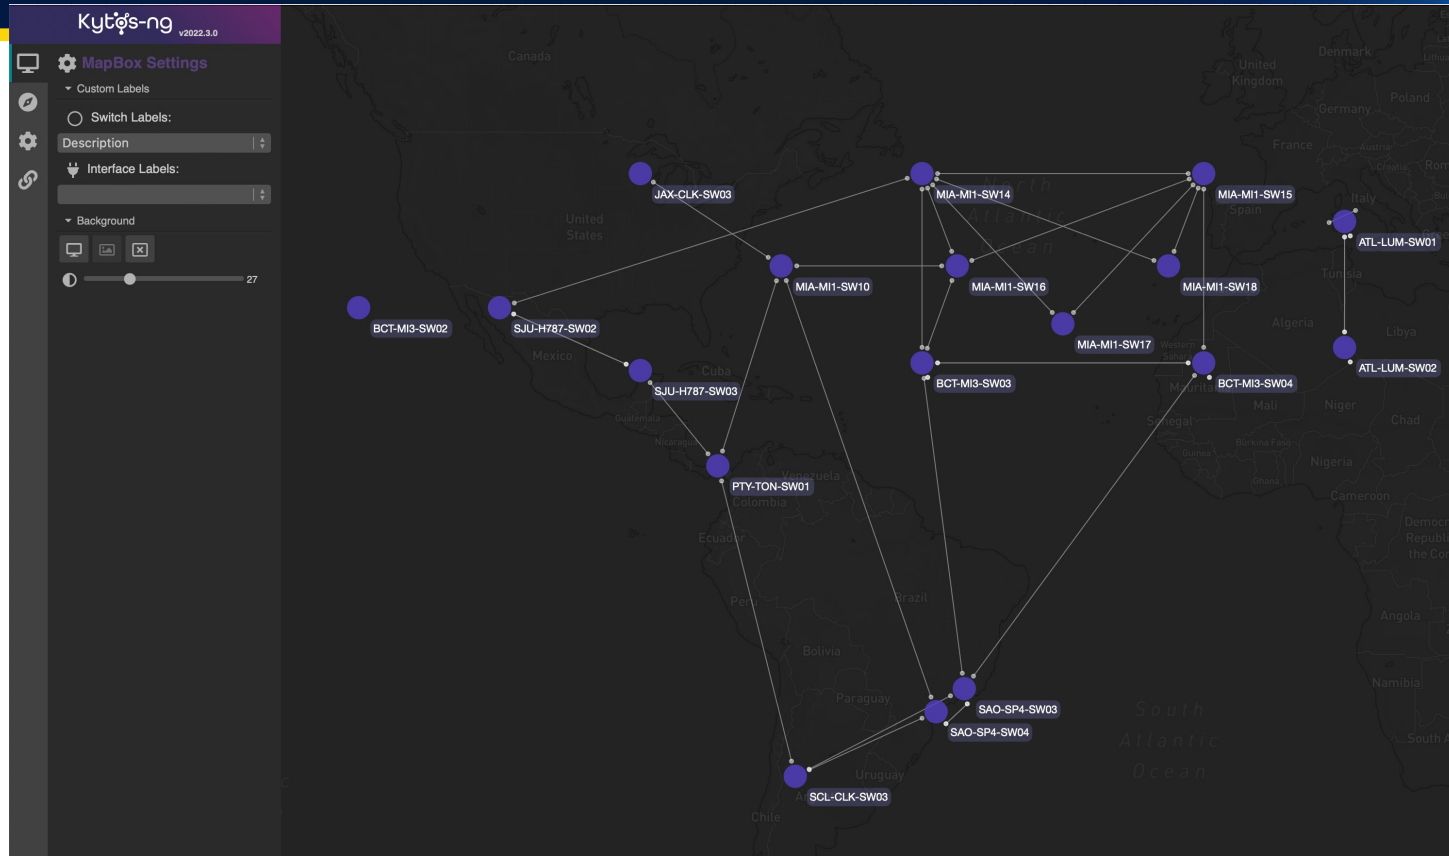
# Control Plane - New Controller: Kytos-ng

- Development focused on the AmLight operation requirements:
  - Simple REST API following OpenAPI 2.0 specs
  - **Pathfinder with support for multiple metrics and restrictions:**
    - # of hops, minimum delay, max bandwidth, ownership, reliability, priority, average bandwidth utilization
  - Supports OpenFlow 1.3+ and gRPC (to support BFRuntime in the future)
  - Fully compatible with NoviFlow OpenFlow experimenter actions
  - Supports Ethernet Private Line (EPL) and Ethernet Virtual Private Line (EVPL)
  - Simple and fast development of new applications
  - Roadmap for 2023/24:
    - Integration with In-band Network Telemetry to add per-packet telemetry
    - Support per-flow Bi-directional Forwarding Detection (BFD)
    - Support for VLAN range

The logo for Kytos-ng, featuring the text "Kytos-ng" in a stylized purple font with a small circular icon containing a dot and a line, resembling a network node or a stylized 'o'.

*Kytos-ng is maintained by FIU and rednesp since May 2021*  
<https://github.com/kytos-ng>

# Control Plane - New Controller: Kytos-ng



# Control Plane - New Controller: Kytos-ng

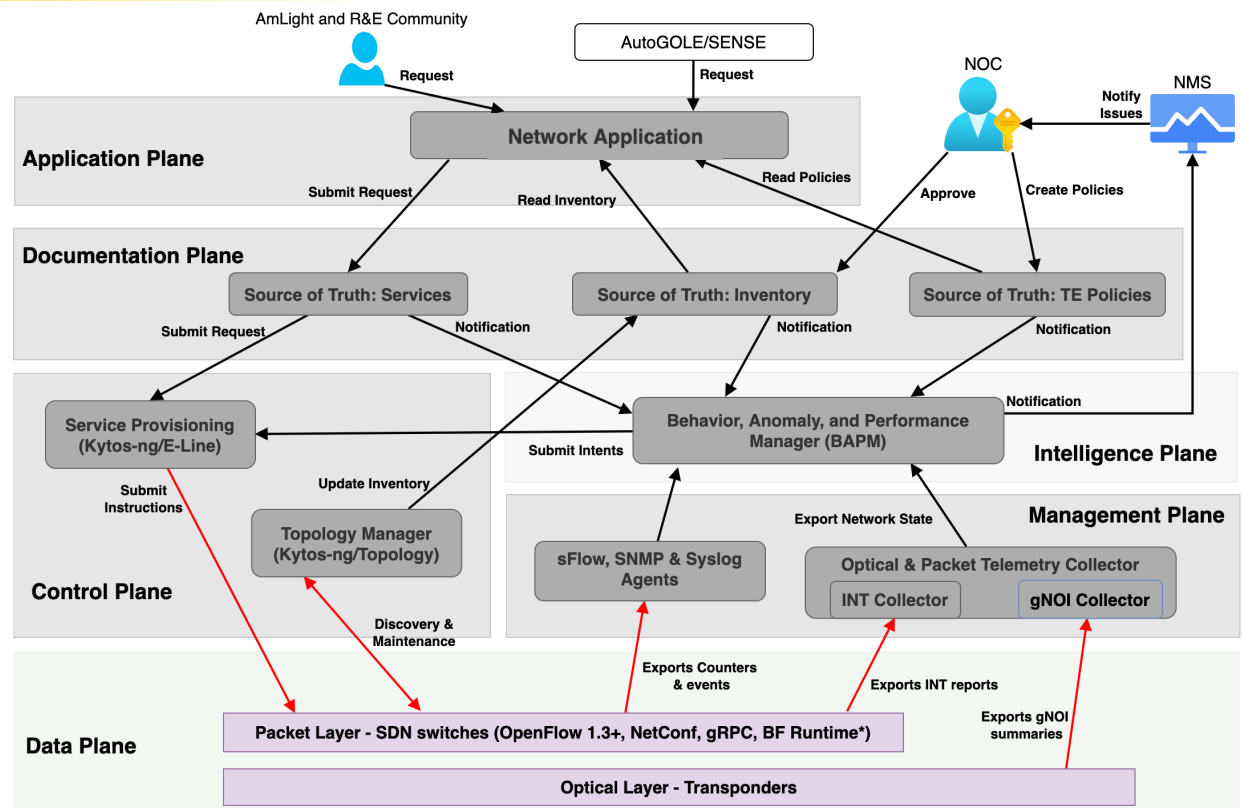
Kytos-ng v2022.3.0

View Connections  
by kytos/mef\_eline

#	Name	Switch A	Port A	Interf. A	VLAN A	Switch Z	Port Z	Interf. Z	VLAN Z	Enabled	Active
1	Vlan_1507_1516_rednesp_amlight_demo_sc23_part1	MIA-MI1-SW14	27	novi_port_27	1507	MIA-MI1-SW14	32	novi_port_32	1507	✓	✓
2	Vlan_1978_ANSP_TENET_part2	SAO-SP4-SW04	1000	novi_lport1000	1978	SAO-SP4-SW04	32	novi_port_32	1978	✓	✓
3	NSI-VLAN-1787-1799	MIA-MI1-SW18	9	novi_port_9	1787	JAX-CLK-SW03	3	novi_port_3	1799	✓	✓
4	Vlan_940_rednesp_academic_pacific_part1	MIA-MI1-SW14	9	novi_port_9	940	PTY-TON-SW01	2	novi_port_2	940	✓	✓
5	Vlan_2011_RNP	MIA-MI1-SW14	9	novi_port_9	2011	SAO-SP4-SW03	25	novi_port_25	2011	✓	✓
6	Vlan_133_FIU_Fabric_I2_part2	MIA-MI1-SW10	3	novi_port_3	133	MIA-MI1-SW17	7	novi_port_7	133	✓	✓
7	Vlan_1623_RNP_Transport_Embrapa-Te	SAO-SP4-SW04	1000	novi_lport1000	1623	SAO-SP4-SW03	25	novi_port_25	1623	✓	✓
8	Vlan_347_346_MIA-Loop_Sp3	MIA-MI1-SW10	1	novi_port_1	347	SAO-SP4-SW03	10	novi_port_10	346	✓	✓
9	Vlan_1505_1513_rednesp_amlight_demo_sc23_part3	SAO-SP4-SW04	32	novi_port_32	1399	SAO-SP4-SW04	18	novi_port_18	1513	✓	✓
10	Vlan_2519_NSI_static_transport_part2	SAO-SP4-SW04	2	novi_port_2	2519	SAO-SP4-SW03	8	novi_port_8	2519	✓	✓
11	Vlan_2501_NSI_static_transport_rnp_part2	SAO-SP4-SW03	30	novi_port_30	2501	SAO-SP4-SW03	1	novi_port_1	2501	✓	✓
12	Vlan_250_RedClara_Shim-RENATA-TekniusColombia	PTY-TON-SW01	318	novi_port_318	250	PTY-TON-SW01	1000	novi_lport1000	250	✓	✓
13	Vlan_2516_NSI_static_transport_rnp_part2	SAO-SP4-SW03	30	novi_port_30	2516	SAO-SP4-SW03	1	novi_port_1	2516	✓	✓
14	Vlan_1508_1518_rednesp_amlight_demo_sc23_part4	SAO-SP4-SW04	3	novi_port_3	1508	SAO-SP4-SW04	1000	novi_lport1000	1518	✓	✓
15	NSI-VLAN-1791-1791	MIA-MI1-SW10	3	novi_port_3	1791	MIA-MI1-SW18	9	novi_port_9	1791	✓	✓

# Data Plane - Forwarding and Telemetry

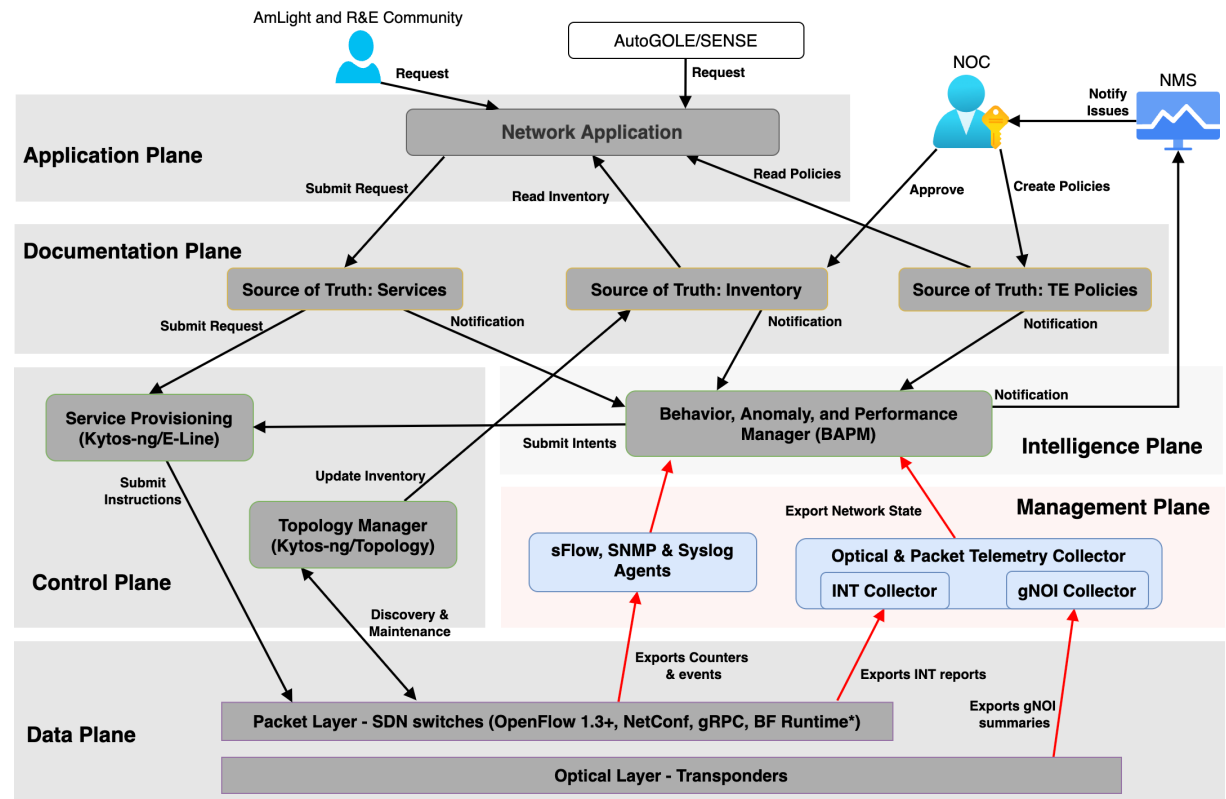
- Second major transformation compared to the previous AmLight-ExP project
- Legacy switches replaced by Edgecore switches and NoviFlow's Noviware
- Support for 100Gbps and 400Gbps
- Support for OpenFlow 1.3+1.4 & BFRuntime
- Supports In-band Network Telemetry (INT) for per-packet telemetry
- From the optical layer, we collect RX power, Q-Factor, and other metrics for proactive monitoring





# Management Plane


- In-band Network Telemetry (INT) exports telemetry reports per-second per-packet
  - Instantaneous bandwidth utilization
  - Instantaneous queue/buffer occupancy
  - Instantaneous hop and flow delays
  - Per-packet path trace
  - Microburst-detection
- Juniper Telemetry Interface (JTI) exports telemetry reports from Juniper MX204 routers:
  - Every 2 seconds for interface counters
  - Every 1 second for device's sensors
- Syslog, SNMP, and sFlow for legacy monitoring




# Management Plane: Network Utilization

- 5 data transfers/bursts of 40-50Gbps for 5 seconds.
- Top: INT metadata exported in real time, per packet
- Bottom: SNMP get running as fast as supported by the switch: 14 seconds.


Interface 11 Utilization - Monitored using In-band Network Telemetry




Interface 11 Utilization - Monitored by SNMP every 15 seconds



# Management Plane: Queue Occupancy



Average Buffer Utilization



Under-Congestion Buffers

# Intelligence Plane: Network Optimization

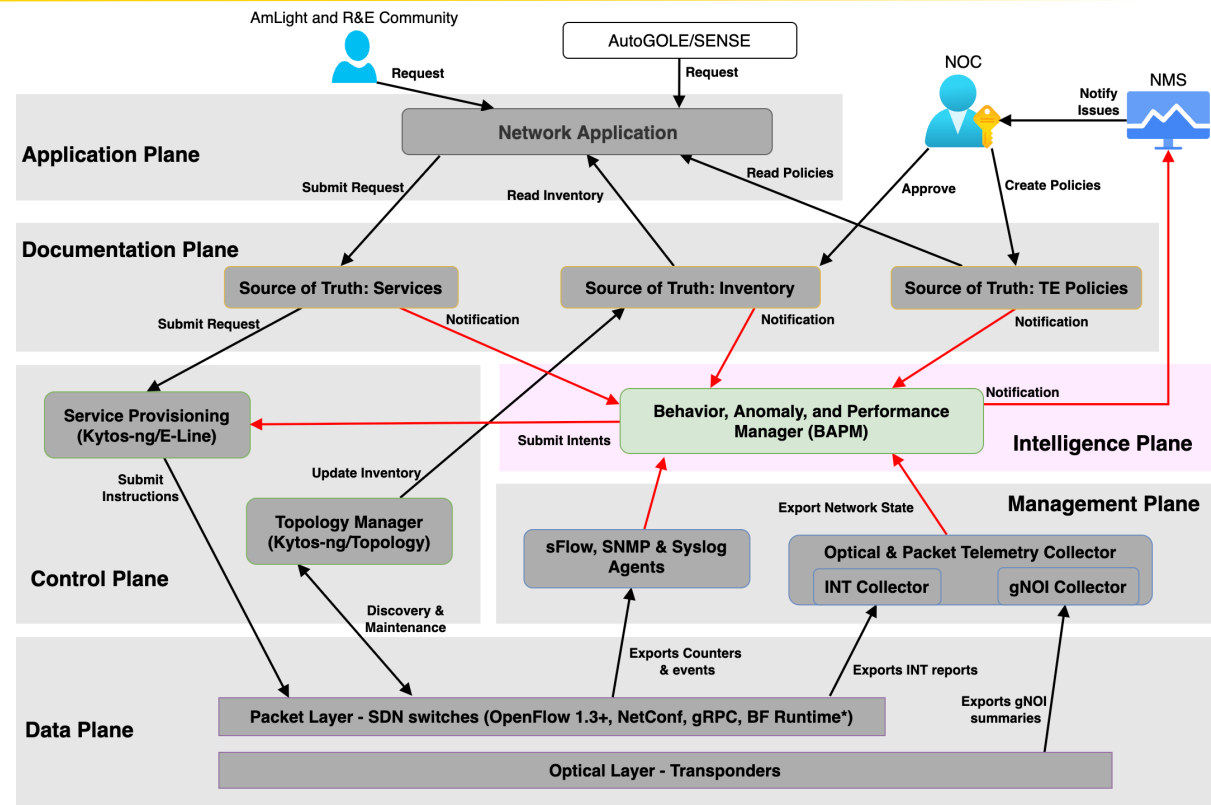
1. Gets inventory, policies, and services from the Documentation Plane
2. Gets telemetry reports from the Management Plane
3. Profiles AmLight's traffic every 100-500ms
  - Discovers performance issues and traffic anomalies
4. Makes suggestions to the Control Plane
  - Steer traffic, Load balance services, rate-limit, etc.

Creates a closed-loop for **network optimization**

- Goal is **sub-second reaction** and debugging
- Behavior, Anomaly, and Performance Manager (BAPM) is the component responsible for the intelligence

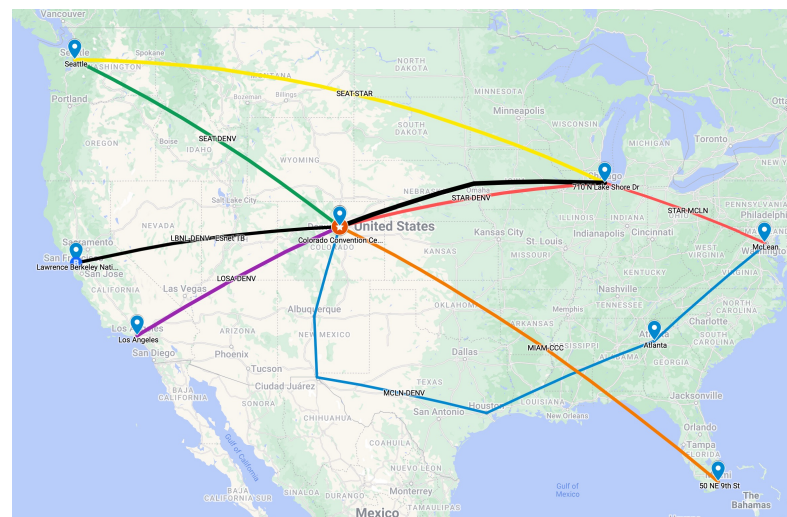

Example of policies:

- 80+% BW utilization  $\geq 2s$
- 50+% [Queue Occupancy]  $\geq 2s$
- Number of path changes  $\geq 5$  in 2h



# Plans and Demos for SC23

- Evaluating AmLight's new capabilities at scale:
  - Leveraging native 400G ports to support research
  - New 100G links and programmable switches
  - New network control with Kytos-ng and its pathfinder application
  - At-scale INT monitoring and visibility
- Providing a platform for connectors to showcase their new infrastructure:
  - rednisp's new 100G backbone with 4x100G connections to AmLight
  - HPC in South Africa leveraging the AmLight's 100G link to Africa (top figure)
- Supporting diverse experiments and NREs with the amazing support from Caltech and SCInet (bottom figure):
  - Global P4 Lab
  - AutoGOLE/SENSE: End-to-End Network Services and Workflow Integration
  - High performance networking with São Paulo Backbone SP connecting 8 universities
  - PolKA routing approach to support traffic engineering for data-intensive science







**AmLight<sub>EXP</sub>**  
*Americas Lightpaths Express & Protect*

Questions? Thank you!

**AmLight 2.0: Flexible control, deep visibility, and programmability @ Tbps!**

Jeronimo Bezerra - FIU/AmLight